ZERO RETRIES

JOURNAL OF THE

NORTHWEST AMATEUR PACKET RADIO ASSN.

Vol. 1 No. 6 September/October 1985

Some Preliminary Observations on HF Forwarding

Having had a decent HF rig for about a week as of this writing, I have been attempting to contact the other HF GateWay operators in the USA. My beam is generally pointed at New England but I have tried other orientations. Those GateWays that I have contacted are indicated with a "*" in the following table.

HF GateWays

	Call	Name	Location	Beam Point	FWD Min	Packet Assoc.
	N1DL	---	Sudbury, MA	78		---
	KC2TN	---	Atco, NJ	84		---
*	K3VPZ	Chuck	Granite, MD	87	3Ø	MAPRC
	K4AHO	---	Orlando, FL	1Ø8		---
	N4CI	---	Atlanta, GA	1Ø2		---
*	K4EID	Bob	Knoxville, TN	98	23	---
*	KF4JF	Larry	Hahira, GA	1Ø2		---
	K4NTA	---	Stuart, FL	1Ø8		---
*	WA4SZK	Don	Florence, SC	98		---
	KI4XO	---	Marietta, GA	1Ø2		---
*	W5XO	---	Gause, TX	121		---
	KD6SQ	---	Rancho Cucamonga, CA	17Ø		---
	WB7DCH	Tom	Enumclaw, WA	---	54	NAPRA
	K7PYK	Wes	Scottsdale, AZ	147		---
*	WB8WGA	Bob	Gahanna, OH	86		None
*	KIØQ	Dave	Ames, IA	95		---
*	WØRLI	Hank	Westford, MA	78	12	NEPRA

In general, HF GateWays operate with DIGIPEAT OFF, a low retry count, long packets (128 bytes), and try to keep four packets in flight at any one time (MAXFRAME 4). The reason for this is simple - WHEN THE PATH IS THERE - USE IT! But even using these groundrules, HF packet forwarding is S L O W !

How long does forwarding take? My present FWD.TNC file attempts to connect with 15 or 16 different MailBoxes and GateWays (If I have traffic for them). While trying to contact all the HF GateWays in the country, I had traffic for ALL of them. What this means is that if the band is flat, I have to retry out (6 retries) 15 or 16 connect requests - This is the FAST case! When the band is open,

and my polite little TNC waits for a moment of silence before transmitting, an unsuccessful connect attempt times out in four minutes. If you connect and you time out or retry out before your traffic is passed, each FWD.TNC entry can take forever! Let's see now, open band, 16 connect attempts at 4 minutes per attempt... That's 64 minutes which leaves minus 4 minutes out of each hour available for users on 145.01. Well, that solves the traffic problem!

What are some of the other problems? Well, for one thing, non-MailBox HF users are a problem. Why should some operator settle for a QSO with his buddy across the country when he could be dumping the entire contents of the disc drives on some remote MailBox? The MailBoxes themselves, during automatic forwarding attempts, cause another problem. The remaining problems are caused by foreign amateurs, CW, RTTY, and the other digital modes. (Have I failed to offend anyone?)

Seriously though, what are some of the steps that can be taken to let us use HF as a temporary "backbone" for cross country linking? Harold Price, NK6K, divides packet radio enthusiasts into two groups: those interested in creating the networks and those interested in using the networks. Obviously, neither group can exist without the other; packet is of little use without networks and networking is no fun without users.

Remember: feel free to copy any or all of this newsletter and share it with others. We ask that any reprinted material will credit both the author and NAPRA's own: "Zero-Retries".

HAPPY PACKETS!

First of all, along with many others, I have never believed that individual users and network trunks can coexist. My exposure to HF operation has been the proof of the pudding. The first option seems to be to separate the GateWays and the HF packet operators. This might be done by moving the GateWays to another frequency on 20 meters and accepting connects from only known GateWays. This won't keep HF packeteers off of the frequency, but will surely discourage its use. Another possibility might be to move GateWay operations to 30 meters. Unfortunately, most of the newer rigs now have 30 meters, and additionally, directional antennas become a problem. Petitioning the FCC for special authorization to allow GateWays to operate 1200 baud with a 1KHz shift on 30 meters may be a possibility. None of these suggestions precludes the use of the existing virtual circuit capability which allows a VHF user in one geographic location to connect to a VHF user in another geographic location via HF GateWays; in fact, the current version of the W0RLI software will remember the correct VHF path to use on the other end. As a Datagram proponent, though, this strikes me as tying up an awful lot of resources while two guys type at each other with their elbows.

Another step which must be taken is the coordination of forwarding times among the HF GateWay operators. At present, forwarding can occur only once per hour. If the GateWay is connected to a user on either the VHF or the HF side when the magic minute occurs, the forwarding gets postponed until the next hour. Without some coordination among the GateWay operators, HF forwarding at one

GateWay can be delayed by another GateWay forwarding traffic to the first GateWay. I plan to undertake this coordination effort. I also intend to limit the number of GateWays to which I forward to three or four. In order for this to work, forwarding agreements will have to be reached with the other GateWay operators.

Currently, information in the network flows only in the forward direction; that is, I have to connect to the remote GateWay to pass my message traffic to him but cannot "reverse the channel" to accept message traffic from him. He must reconnect to me at a later time to pass his traffic to me. This seems to be a real waste of channel capacity. If the path exists and the connection has been established it seems that traffic should flow in both directions.

One of the major drawbacks to HF GateWay operation is that the MailBox is unavailable on the VHF side during HF forwarding or when in use by someone on the HF side. This tends to leave the local users with the feeling that they are supporting a pretty unreliable MailBox. A couple of solutions suggest themselves: the MailBox/GateWay software could be "dual-ported" in order to allow simultaneous operation on both sides, or alternately, the MailBox/GateWay shouldn't be the main local VHF MailBox. In the absence of a rush by volunteers to whomp together additional Xerox 820 MailBox systems, perhaps we'd best examine the first possibility. In concept, the system could be dual-ported. In practice, there are two major difficulties: the fact that CP/M will not adequately support multiprocessing, and, even in a multiprocessing environment (ie MP/M), the processing and I/O capabilities of the X-820 will leave much to be desired. The second possibility really seems the most attractive. This frees up the HF GateWay to devote the majority of its time to handling traffic in and out of an area, while leaving the local VHF MailBox available for local users. The only times the systems would have to talk to each other is while passing traffic between the GateWay and the MailBox.

Now that we have a local HF GateWay, how do we go about using it? The key to its use is knowing at which local MailBox or MailBox/GateWay the person you wish to contact receives his mail. You use the call of that MailBox in the "@" portion of the address. For example, I recently wanted to communicate with Tom Clark, W3IWI, near Washington DC. From the following table of known MailBoxes, I know that Tom Clark runs the W3IWI Mailbox. I can also safely assume that he receives his mail there. What I can't tell from the table is that the nearest HF GateWay is K3VPZ and that K3VPZ forwards (via some means - probably VHF) to W3IWI. In this case, I was told by W0RLI that K3VPZ forwards to W3IWI. As the operator of the WB7DCH MailBox/GateWay, I entered the call of the W3IWI MailBox in the FWD.TNC file under K3VPZ. Now, as long as I continue to forward to K3VPZ, mail addressed to anyone who receives his mail at W3IWI will be automatically forwarded to K3VPZ who will handle it from there. Another example: suppose you wish to send a message or an article to Jeff Ward, K8KA, who edits the ARRL Gateway newsletter. From the table, you know that W1AW, at ARRL Headquarters, is a MailBox. You simply address your message to K8KA @ W1AW and, as if by magic, it will arrive at the W1AW MailBox. How does

MailBoxes

Call	Location	Call	Location
W1AW-4	Newington, CT	AK3P	Hummelstown, PA
K1BC	Lexington, MA	KS3Q	Baltimore, MD
KE1G-1	Goffstown, NH	WA3QFN	Margate, FL
N1DKF	Cranston, RI	K3RLI	Wilkes Barre, PA
WB1DSW	E. Kingston, NH	NK4K	Miami, FL
WA1LRL-1	Beverly, MA	K4TKU	Miami, FL
WB2MNF	Medford, NJ	WB4ZNW	Moultrie, GA
AI2Q	Freeport, NY	WB5VUL-1	Palo Alto, CA
WA2RKN-2	Hyde Park, NY	WA5MWD	Garland, TX
WB2RVX	Pennsauken, NJ	W6CUS-1	Richmond, CA
WA2SNA-1	Wayne, NJ	WAØJFS-1	Des Moines, IA
WA2VKH	Carlstadt, NJ	KCØOJ	Battle Creek, IA
N2WX-1	Palm Bay, FL	NAØP	Ames, IA
KC3BN	Harrisburg, PA	NAØS	Ames, IA
WB3EYB	Harrisburg, PA	AIØZ	Roland, IA
K3GYS	Atco, NJ	VE3FXI	Ottawa, ON
W3IWI	Clarksville, MD		

---

this happen? My FWD.TNC file lists W1AW as one of the MailBoxes served by W0RLI. Any mail which specifies "@ W1AW" is automatically forwarded to W0RLI, who forwards it on from there.

What can you do to help? There are six X-820 boards in the area that I know about and probably a few others whose owners are keeping an even lower profile. At present there is a tremendous amount of talk and enthusiasm about MailBoxes, dual- and quad-ported digipeaters, link management units, and expanding our local area network to all the points of the compass. We've had this same tremendous amount of talk and enthusiasm for the last year. That's all we've had - talk and enthusiasm. Those of us who have accomplished anything can be counted on the fingers of one hand. I might also add that these same doers are getting pretty burned out, both emotionally and financially. If you are reading this, you have probably invested in a TNC. What will your investment be worth if the network never gets completed? How useful would your 1-watt handheld be without repeaters? Get involved! Attend the meetings; let your voice be heard; be willing to assume the responsibility for getting something done!

WB7DCH 85/8/30

BETA TESTING THE NEW PK-64

This month AEA will introduce the newest and latest product to enter the packet radio market. Dubbed the PK-64, this product is obviously aimed at owners of Commodore computers. But Commodore is in for a surprise! Hams who are interested in packet that still don't own computers now have a good reason to purchase the C-64.

The PK-64 will do nearly anything that TAPR's TNC-2 will do. Infact, it uses the same commands. Unlike TNC-1 and -2, TAPR clones and Kantronics TNC's, the PK-64 uses a menu page to set all but a few of the many parameters. Another menu page allows the user to select a variety of options that include loading and saving buffers to disk or cassette, editing buffers, browsing through buffers, the list goes on and on.

Although it was not present in the beta board I have been using, rumors are out that the list goes well beyond packet and that the product will turn your shack into a complete digital station. We'll have to hear the facts from AEA.

Split screen terminal software has been cleverly integrated into the package. Not only does this board turn your C-64 into a TNC, but it provides you with exceptional software for easily operating your packet station. If you hook up a speaker the computer a siren alerts you to the fact that someone has connected to your station. Your computer will call you in from another room. If you aren't around when a call comes in, a message is automatically sent prompting the caller to leave a note that you will read later. A timer provides the security of knowing that your station will not be left "hung-up" by someone who has retried out.

There are 10 user buffers that can be easily loaded from disk or tape and used to send full messages with the touch of a few keys.

A fairly adequate text editor allows you to manipulate data in these buffers before you send them or save them.

Would you like to try HF packet? The PK-64 will switch to HF speeds and tones with the touch of a key on the keyboard. And AEA has come up a with filtering system that makes HF packet easier than ever before. I ran the beta board and my TAPR TNC-1 side by side. The PK-64 was capable of copying packets that the TNC-1 never ever heard!

The innovative work that AEA has put into this new product will shake up the competition. There is no other product like it. Those who scoff at the C-64, calling it a kid's toy, will have to think again when they see it hooked up to the PK-64. Price, performance, and surprises; I give the PK-64 an A+ all the way down the line.

Budd Churchward, WB7FHC

TNC 2 Assembly Update

It took me a total of four hours to assemble the first of three TNC-2's at AEA. Defluxing the PCB took alot of Q-Tips and freon... and things would've been much easier if I had found the assembly errata sheets before I completed the kit, and not after!

It took some additional time to figure out that I had tuned the demodulator to the second harmonic of the desired frequency.

After talking to several other TNC 2 builders on the phone over the past several days, it seems that everybody had problems with damaged DIP sockets. In the first kit, I had one each 16, 28, & 40 DIP socket smashed, one smashed LED, along with a missing 0.1uF by-pass cap, and one severely maimed lithium battery. The battery, luckily, was salvagable. A discussion with Pete Eaton down at TAPR brings confirmation that all of the above components will be packaged differently in future kits.

73 de K Brett 7 Graham, KB7G

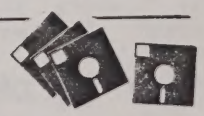

cmd:PREZ SEZ ...

What is the life blood of an organization? It is the members getting involved. We on the Board of Directors try to think of your wants and desires, but have a hard time of it since we rarely hear from you. We desperately need to know what you think! Drop us a line from time to time. Any feedback will be appreciated.

The time for elections is just around the corner. At the September meeting at the Radio Club of Tacoma clubhouse we will be accepting nominations for all the officer positions. At the October meeting at the BEARS clubhouse at Renton Air Field we

will accept further nominations with elections to follow. Please come and participate.

From now on the technical committee will have a meeting at 9 am on the second Saturday of the month - the morning of the regular 1 pm club meeting. It will be at the same location as the regular club meeting. Listen to the Thursday night ragchew (145.33 RPTR) if its location will be changed. The purpose of this meeting is to exchange ideas with one another on projects in the planning stage or under construction. All that is required for attendance is curiosity.

I wish to thank all those who helped in the fine showing at the Tacoma Hamfair. Their hard work and many hours in preparation and performance were very evident. The information table was able to show packet radio in operation and answer many questions of would be packeteers. Likewise, the WB7DCH MailBox was an excellent demonstration of some of the features and capabilities of packet radio. The packet radio presentation by the guest speaker Harold Price, NK6K, was well attended. Again, many thanks to all those who made it happen.

Dennis Goodwin, KB7DZ
President of NAPRA

TAPR TNC ON 300 BAUD - 200 SHIFT

TAPR advises the following if the TNC-1 will not receive when using 1600-1800 Hz shift at 300 baud.

Simply transmit once at 300 baud like:

```
cmd: C W1ABC
^C
cmd: D
```

Once the TNC-1 has transmitted it will then receive fine on HF.

I know of three units that require the above to make operational on HF.

Also, it is normal, on power up at 300 baud to get the message: "HDLC CAN'T INIT".

Gervie, W7FEN

DOWNLOADED:

PSR QUARTERLY JULY 1985

THE TWO PORT DIGIPEATER

Jon Bloom, KE3Z

Recently, I have heard from several people who seem to be confused by the mutli-port digipeater concept. What I hope to do with this article is to reduce the confusion level. I'll first explain the nature of the problem that prompted the development of the multi-port digipeater. Next, I'll explain what the digipeater does. Then I hope to show exactly how the multi-port digipeater solves, or at least reduces, the problem. Finally, I'll look at some interesting applications of the multi-port digipeater.

THE PROBLEM

Anyone who is more than casually interested in amateur packet radio will recognize the problem. I'll describe it from the perspective of a user of the northern part of EASTNET, although the same problem is found to a greater or lesser extent in all areas of the United States.

To understand the nature of the problem, we have to understand an entity known as a "collision" (please forgive me if I'm treading old ground for some of you). A collision occurs when a receiving station hears two transmitting stations at the same time. When this happens, the receiving station cannot properly receive the transmitted frame(s) of both transmitting stations. The nature ot FM, which is the modulation method used on VHF packet, is such that the receiver MAY recive one of the signals if it is considerably stronger than the other, but often neither signal is received. For the purposes of our discussion, I will assume that neither transmitted frame is properly received.

When I wish to connect to another station using the local digipeater, there are three points at which a collision can occur. When my TNC transmits the initial SABM (connect request) frame, it may collide with another transmission at the receiver of the local digipeater. If that occurs, the digipeater will not retransmit the frame. My TNC will wait for a period of time and then try again. It I manage to transmit my SABM to the local repeater without a collision, the repeater will then retransmit the frame. Again, there is the possibility of a collision. This is usually not a problem because the local digipeater is by far the strongest signal on frequency at the receiving station. When the staion I'm calling responds to my SABM with a UA (acknowledge) frame, the same potential for collisions occurs. This time, the receivers involved are, once again, the digipeater, and finally my receiver. Because of the ARQ nature of packet, retries will usually "solve" the collison problem eventually, although a significant delay occurs because of the time between retries. Unfortunately, this is the simple case. Many packet connections occur over multiple-digipeater paths. In this case, there is the potential for collisions at the receiver of each digipeater in the path. As pointed out by W3IWI, the probability of reception at the destination station decreases exponentially as the number of digipeaters increases. For example, if the probablility of reception for each digipeater in a five-digipeater path is 90% (an optimistic value), the probability ot reception at the end point of the path is 59% (.9^5). This lack of reliability results in significant network delays.

Where do the frames that collide with mine come from? Most of them are NOT from other local users. Rather, they come from other digipeaters. As an example, the local repeater here in central Connecticut is W1AW-5, in South Windsor. This machine can hear four other wide-coverage digipeaters: KG1O-9, Mt. Ninham, NY; KY1H, Peru, MA; WA1TLN-1, Mt. Ascutney, VT; and, WA1IXU, Collinsville, CT. During the "prime time" hours of the evening, at least one of these digipeaters is transmitting at any time. This makes it difficult to get ANY data through the network.

The problem, therefore, comes down to one basic fact: with all of the digipeaters in the network on the same frequency, collisions are impossible to avoid.

THE MULTI-PORT DIGIPEATER

The multi-port digipeater was developed as a temporary solution to the problem. I say temporary because, if justice triumphs, it will be replaced with true network-layer machines in the not-too-distant future. The major effect of the multi-port digipeater will be to get the network digipeaters, and hence their local users, onto different frequencies.

The multi-port digipeater has two or more "ports" which are separate digipeater nodes. Typically, each port will be on a different frequency -- perhaps even on a different band. Each port is capable of acting as a stand-alone digipeater. That is, it can receive and transmit frames in the normal digipeater manner. But the multi-port digipeater can do more, as it can route frames received on one port out a different one. How does it do this?

Within the EPROM memory of the two-port digipeater are several constants and tables which determine the transmission port for a received frame. The digipeater determines the transmission port by examining the address field of the received frame (routing). If the frame is fully repeated (this digipeater is the last one in the address field of the received frame), the digipeater software first looks through a table of "destination stations" for a match with the address of the destination station in the received frame. If it finds a match, it retransmits the frame via the port indicated in the table.

If the received frame is NOT fully repeated (there is a digipeater in the address field after this digipeater), a different address table is scanned, looking for a match with the address of the next digipeater. If a match is found, the digipeater will, once again, retransmit the frame via the port indicated for that address by the table. If no match is found, the SSID of the next repeater address is examined. The multi-port digipeater can be configured to send all frames of a specific repeater SSID out a particular port. For example, all frames in which the next repeater address has an SSID of 8 could be routed out port 1. Each port can have its own "default" SSID.

Should explicit routing fail (that is, the address/SSID was not found in the internal tables), the digipeater retransmits the frame on the default retransmission port for the receive port. If no explicit routing is determined the digipeater may retransmit a frame on the same port it was received on. Thus it will "look" like a normal single-port digipeater to that frame: received and retransmitted on the same port, or frequency. In the proposed network, all frames that can't be routed by SSID or table look- up will be transmitted on the local, or 2-meter port.

This is all very confusing and you are probably wondering what purpose it all serves! In the next section we will see by example how the multi- port digipeater uses its routing capabilities to solve the network collision problem.

THE SOLUTION TO THE PROBLEM

Consider the network shown graphically below. Each digipeater is a two-port digipeater. The frequencies used on each port and the tables set up at each digipeater are shown. All of the inter-digipeater links are on a common 220-MHz frequency. Any frame which is not explicity routed is retransmitted on 2-meters.

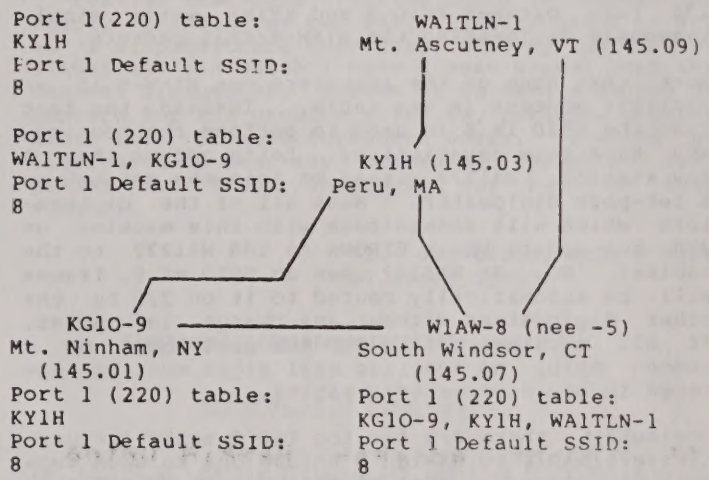

Let's follow a few frames thru the network. WB2KMY in Poughkeepsie, NY wants to connect to KE1G in Goffstown, NH. WB2KMY is on 145.01 MHz and sends his connect request via the path: KG1O-9, KY1H, WA1TLN-1. Upon receiving this frame, KG1O-9 examines it, sees that the next repeater is KY1H and finds KY1H in its port-1 table. It therefore retransmits the frame on 220. KY1H receives it, sees WA1TLN-1 in the address, and retransmits it on 220 again. When WA1TLN-1 receives it, it is a fully repeated frame. Since there is nothing in the TLN-1 destination station address table, the frame is transmitted on the default transmission port, on 2 meters. This default action eliminates the need for a table of local 2-meter users; only the 220-Mhz stations (the "backbone" network) need be known to the digipeater.

Now, KE1G receives the frame and replies via WA1TLN-1,KY1H,KG1O-9. The same routing action occurs on this path, with WA1TLN-1 and KY1H transmitting on 220 and KG1O-9 on 2 meters, where WB2KMY receives it.

Let's say that another activity is occurring at the same time. In this case, KE3Z is reading a long message from the W1AW-4 PBBS, going through W1AW-5. When a frame is transmitted by W1AW-4, W1AW-5 receives it. Since W1AW-5 is the only repeater in the addres field, the frame is fully repeated. Neither the address KE3Z-0 nor the SSID 0 are to be found in the tables of W1AW-5. So, it retransmits the frame on the default, 2-meter port. In other words, all of this activity at

W1AW-5 is occurring on 145.07. NO TRANSMISSION IS MADE ON THE 220 LINKING FREQUENCY. Since none of the other repeaters have a 2-meter frequency ot 145.07, none of the activity at W1AW-5 is heard by the other machines and vice versa.

So what happens when there are several connections active via the 220 linking frequency? Aren't we still going to have collisions? Yes, but far fewer. Remember, any connection which only makes use of a single digipeater uses no part of the available 220 capability. Also, since every station on the 220 frequency is a wide-area digipeater with a good VHF location, most of them can hear one another. This is a situation in which the Carrier- sense/Multiple-access (CSMA) protocol used in packet radio works well. Local users on 2 meters can transmit at the same time a distant digipeater is transmitting on 220 without collisions. In fact, W1AW-8 can digipeat packets on 220 (say between KG1O-9 and KY1H) while simultaneously digipeating the W1AW-4/KE3Z packets.

Note that none of the repeaters has W1AW-8 as an explicit address in its table. Instead, the fact that the SSID is 8 is used to perform the routing. Why have this capabiltity? Let's assume that a new station, WA1??? wishes to join the network as a two-port digipeater. Must all of the digipeaters which will communicate with this machine on 220 now update their EPROMs to add WA1??? to the tables? No. It WA1??? uses an SSID of 8, frames will be automatically routed to it on 220 by the other digipeaters without any change in tables. If all machines with a 220-MHz port agree on a common SSID, no specific call signs must be entered in any digipeater's tables.

Obviously, there are far too few 2-meter frequencies available to assign a unique one to each two-port digipeater. What we end up with is a situation somewhat analogous to that of the voice repeaters. We could, for example, have a machine in New Hampshire on 145.07 because it would not normally hear, or be heard by, W1AW-8. Still, there are far fewer frequencies available to 2-meter packet users than voice machines. Fortunately, packet repeaters can accept a small amount of co-channel interference better than voice machines can.

DEBUNKING THE SSID-TRANSLATION TECHNIQUE

When the multi-port digipeater concept is first considered, there is an intuitive belief that some sort of SSID translation is needed. That is, that the second port of the machine operates with a different SSID than the first port. I was bedazzled by this illusion too, but I eventually discarded it. To briefly explain the concept: When such a machine received a frame on port 0 (the 2-meter port), it would retransmit it on port 0 if the received SSID of its own address was 1. It would retransmit it on port 1 (the 220 port) if the SSID was 2. When it retransmitted the frame on 220, it would have translated its SSID in the frame to 1. This is so that any reply frames generated by the far endpoint station would have been generated with an SSID of 1, forcing the retransmission of the reply frame on port 0. (If this seems confusing to you it's because it's confusing.)

My main problem with this concept is that it requires the user to have an intimate knowledge of the ports used at each digipeater in the path. You have to specify the retransmission port of each digipeater by selecting the proper SSID. This is not TOO awful as long as all of the digipeaters are two-port machines; all of the digipeaters are -2 (220 backbone) except the last one, which is -1 (for the 2-meter destination). If, however, one or more of the digipeaters is a three- or four-port machine, or if several repeaters are operating under the same call sign, it can be very hard to get the path right (we will discuss the function of digipeaters with more than two ports shortly). What the SSID-translation technique accomplishes is to remove the routing intelligence from the digipeater and force the user to supply it. This seems to me to be a giant step backwards.

The other major failing of the SSID-translation scheme is that it does not inhibit direct user access of the linking frequencies. If users start showing up on the "trunk" lines, we will rapidly be back just where we are now, with large numbers of collisions and a network that grinds to a halt. The SSID-translation machine doesn't protect against this, as it makes no difference which port you access. With the intelligent routing machine, a user accessing the machine via the network port never gets a reply from a called station. This is because when the reply frame is received back by the digipeater, it is fully repeated and goes out the 2-meter port by default.

THE TRUE MULTI-PORT DIGIPEATER

So far, we have shown only the two-port machine in operation. In the simplified network shown above, two-port machines handled everything nicely. In reality, a nagging problem crops up as the network is expanded. Let's say that the 220 link is expanded down across New Jersey from KG1O-9. Since the New England machines can't hope to hear the New Jersey machines, the potential for collisions between New Jersey frames and New England frames at the input ot KG1O-9 is very high. This is not to say that the network wouldn't be much improved over the present state, just that the same old problem is starting to rear its ugly head again. One possible solution is to have the link between KG1O-9 and the New Jersey machines on yet another linking frequency -- one which would not interfere with the 220 frequency used in New England. This frequency could be on 220 or perhaps 450, which would ease the implementation at KG1O-9. At the southern end of New Jersey, where linking to the MD/DC/VA/PA areas occurs, the 220 frequency used in New England could once again be pressed into service. The three-port machines used at these "choke points" where much of the network traffic is handled would, of course, be key players in the network.

WB2KMY-1 digipeater is on Mt. Beacon, NY. It is very close to KG1O-9, but is isolated from much ot the New England and New Jersey traffic by mountainous terrain. If it links to KG1O-9 on one of the two 220 linking frequencies, collisions will occur. Therefore, it would be ideal if the link between these two machines were on yet a fourth frequency. With the terrain and distances involved, 1296 MHz or even 10 GHz could be used. KG1O-9 would then be a four-port machine and WB2KMY-1 would need three ports.

Recently, I have heard some naysayers deriding this scheme because "there are too many frequencies involved" or some such specious argument. What these scoffers fail to realize is that even when we have true network-layer machines, the same frequency/topology problems will persist. Indeed,

one of the motives behind the multi-port digipeater project is to get the RF media in place, ready for the next generation of the network: level 3.

NETWORK TRANSMISSION SPEEDS

I haven't mentioned the subject of transmission speeds on the inter-digipeater links before this because it isn't really germane to the basic concept of the multi-port digipeater; the advantages explained above would acrue to the network even of the links are still operated at 1200 bauds. Those of you who follow packet developments closely know that a revolutionary (in amateur terms, at least) new modem has appeared from the lab of K9NG. This modem is capable of operating at speeds greater than 9600 bit/s, although 9600 appears to be the developing standard. I anticipate that no inter-digipeater link will operate at below that speed.

The increase to 9600 bit/s provides an obvious advantage: eight times as much traffic can be carried by a single link. What may not be so obvious, at first glance, is the comcomitant benefit: lower delay. It should take less time to send a frame and get a reply through 8 digipeaters (4 1200-baud transmissions and 12 9600-baud transmissions) than it does through two digipeaters (6 1200-baud transmissions) now. Combine that with the infrequency of collisions and you have a network that really flys! Note also that the speed limitation of the K9NG modem is that of the bandwidth of the RF system. There is no reason why much higher speeds could not be attained with specially-designed radios.

There is also no reason why a 9600-baud local port could not be implemented. For instance, a digipeater in the network could have a 9600- baud network port, a 1200-baud "entry-level user" port and a 9600-baud user port for the more ambitious packeteer. The current 2-port digipeater, which is based on a Xerox 820, will support a 1200/9600, 300/1200 or most any combination of two standard speeds.

OTHER USES OF THE MULTI-PORT DIGIPEATER

This isn't really part of the network discussion, but it hasn't been published anywhere else to my knowledge so I thought I would include it.

The flexibility of the routing scheme employed by the multi-port digipeater gives rise to several other possible uses. One of the more interesting is its use as an HF/VHF gateway. This is done by "cross channel" routing, in which everything received on port 0 is retransmitted out port 1 and vice versa. (Naturally, by "everything" I mean all frames addressed to go through the digipeater!) With port 0 on 2 meters and port 1 on HF, you have a simple, cheap gateway which can be used by multiple users simultaneously. Right now, the only operational HF/VHF gateways of which I am aware make use of the WORLI MailBox software. Such gateways can only be used by one user at a time and require the use of two TAPR TNCs or clones thereof.

Another possibility, albeit one that doesn't excite me, is use as a "closed" packet repeater. By having the addresses of the authorized users in the ROM and "grounding" (throwing away) any received frame from other users, the repeater would be effectively closed. Grounding frames which are not explicitly routed by the tables is accomplised by setting the default retransmission port to route frames to a null I/O driver. Of course, the EPROM would have to be updated to add a user. This would be a pain, but anyone who runs a closed repeater deserves it. It also requires that someone write the null driver or modify the existing software slightly.

Although a multi-port digipeater with more than two ports has not yet been implemented, the software is designed to support it. What needs to be done is to add more HDLC ports to the Xerox 820, or run the software on a different machine with more ports built in. The FAD board from TAPR is one possible solution to adding ports to the 820, although the software is not set up to drive an 8530, the chip used on this board. I hope to hear shortly from some people who have been working on an 8530 driver and will make that available after testing it.

FINAL WORDS

I realize that this treatise could be viewed as simply grinding my own axe. I have tried to act as a dispassionate observer, but my particular situation, in which I have a weak signal into the nearest digipeater, forces me to have a degree of concern for the problem which is, perhaps, greater that that of more fortunate users. I would be interested to hear any comments, thoughts, complaints or Bronx cheers you may have on this subject.

73, Jon Bloom, KE3Z @ W1AW
17 June, 1985

##########

NAPRA Board of Directors:

The following people serve us all and have contributed a lot to make NAPRA work. These are the elected officers of our association. Others have also worked hard. Support NAPRA.

KB7DZ	Dennis Goodwin	President
N7DTI	John Gates	Executive Vice President
KD7UW	John D. Hays	Technical Vice President
WD7FHC	Budd Churchward	Secretary/Treasure
		Editor: Zero Retries
WB7OCH	Tom Hogan	Director at Large

##########

NAPRA Scheduled Meetings:

REGULAR, GENERAL MEETINGS ARE HELD ON THE SECOND SATURDAY OF EACH MONTH AT 1 PM. MEETING PLACE ALTERNATES BETWEEN THE RADIO CLUB OF TACOMA CLUB HOUSE IN TACOMA AND THE BOEING BEARS' CLUB HOUSE IN RENTON (CHECK THE MAILBOX OR THE THURSDAY 9 PM NET ON 145.33 TO FIND OUT WHERE).

BOARD MEETINGS (OPEN TO ALL) ARE HELD EVERY THIRD WEDNESDAY OF THE MONTH AT THE OMNI RESTAURANT IN FEDERAL WAY.

##########

LINKING ON 145.01

To those who did not attend the August BOARD MEETING:

It has been decided to go ahead with plans for a 145.010 Mhz. BACKBONE from Vancouver, BC to Vancouver, WA ... Also from this BACKBONE to Ellensburgh, WA and Spokane, WA.

The key word here is BACKBONE ... Not closely placed LANs ... 5.01 would in this context become a circuit for INTER-AREA communications with new LOCAL Area Networks [LANs] coming on line on 5.05, 5.07, 5.09 and SIMPLEX operations moving to other frequencies including 144.99 Mhz.

We have been waiting for almost a year here in Seattle for a 2-4 Port Digipeater ... and have not attempted out of area networking except on HF thru the DCH mailbox. We need to move beyond our present borders with INTER-AREA communications in order for NAPRA and Packet Radio to grow in the NORTHWEST.

The bleak outlook of some for the effectiveness of this network is acknowleged and understood .... and it is ONLY A TEMPORARY solution to a complex problem. The Technical Committee and Board of Directors recognize this ... however, we want Packet Radio to GROW!

The situation is not as bad as has been outlined as the interconnecting digipeaters from ANK-5 to Vancouver, WA will have a lower altitude and are in areas with fairly small PACKET communities... We would hope that people would learn to USE the BACKBONE for INTER-AREA and BBS access [plus as the Calling Freq.] and move their other traffic to the other LANs and SIMPLEX frequencies.

Our plan is to rapidly move over to a system where many local area networks operating on 2M and 220Mhz. would be interconnected using multi-port digipeaters on an UHF frequency. So if any of you have a MULTI-PORT DIGI sit ting around the shack please contact us ASAP.

"HE WHO WAITS ... GETS LEFT BEHIND" - Anonymous

DE John KD7UW - Tech VP NAPRA

Editor's note: This announcement has generated a great deal of controversy. Subsequent meetings of the Board and the technical committee have resulted in very heated discussions on the merits of John's proposal and solutions to the link problem that envolve the use of more frequencies. Just resently another digipeater, WB7TGF-1, has come on the air on 145.01 from the Olympia area. We are told that they are prepared to QSY to 145.09 as soon as a linking scheme can be implemented. Two 820 bridges are currently under construction and a cadre of several NAPRA members are anxious to show us that there is a better way.

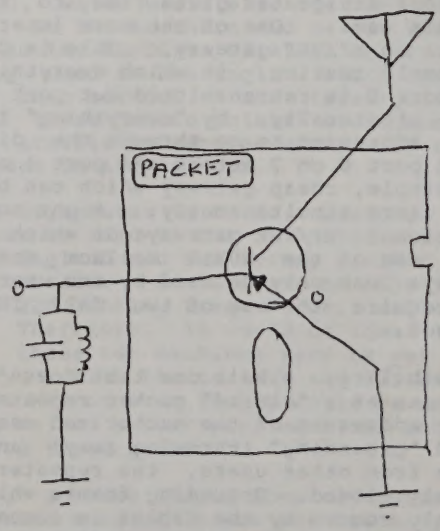

Cal '85

BEGINNER'S GUIDE TO COMMON MICROCOMPUTER TERMS

IBM Compatible: Not IBM compatible.

Fully IBM-Compatible: Somewhat IBM-compatible, but won't run BASIC programs.

100% IBM-Compatible: Compatible with most available hardware and software, but not with the blockbusters IBM always introduces the day after tomorrow.

Lap-Top: Smaller and lighter than the average secretary.

Portable: Smaller and lighter than the average refrigerator.

Transportable: Neither chained to a wall nor attached to an alarm system.

Hard Disk: A device that allows naive users to delete vast quantities of data with simple mnemonic commands.

Mouse: A peripheral originally christened "vermiform appendix" because of its functional resemblance, renamed for its appropriateness as a cat toy.

Printer: An electromechanical paper-shredding device.

Modem: A peripheral used in the unsuccessful attempt to get two computers to communicate with each other.

Network: An electronic means of allowing more than one person at a time to corrupt, trash, or otherwise cause permanent damage to useful information.

Documentation: A perplexing linen-bound accessory resorted to only in situations of dire need when friends and dealers are unavailable, usually employed solely as a decorative bookend.

Version 1.0: Buggier than Maine in June; eats data.

Version 1.1: Eats data only occasionally, upgrade free to avoid litigation by disgruntled users of version 1.0.

Version 2.0: The version originally planned as the first release (except for a couple of data-eating bugs that just won't seem to go away), no free upgrades or the company would go bankrupt.

Version 3.0: The revision in the works when the company goes bankrupt.

Now Available: Available any day now.

Available Soon: Available in a year or so.

Available May 1: Version 1.0 may ship to dealers August 1.

Link State Indicator and Audio Alarm

Bill Johnson, WD6FPY, and his brother, Jim, WA6OZJ came up with this clever TNC accesory. It can be used with TAPR Beta and Kit boards and with the Heathkit version of the TNC. It plugs into the parallel port.

The purpose of the device is to give you a variety of information. It does the following:

1. Sounds an audio alarm when someone connects with you.

2. Turns on an LED on connect which remains lit until you reset it. If you are in ctrl S, you will know if someone has connected.

3. Lights up a seven segment display to let you know what your link state is. Numbers indicate the following:

    0 TNC has been reset.
    1 Disconnected.
    2 Connect attempt in progress.
    3 FRMR condition detected.
    4 Disconnect attempt in progress.
    5 Connected.
    . Decimal point lit, converse mode invoked.

The fellows used a 150 ohm resistor in series with the LED. I found that a 470 ohm resistor was required with the LED I used, to cause the LED to stay lit after disconnect. Where did all this neat information come from? See page 6-6 of the TAPR Kit manual.

Bob Hart - WA7HRA

MEMBERSHIP RENEWAL FORM

Name_____ Callsign_____ Class_____
Address_____
City_____ State_____ ZIP_____
Phone_____ (Days) _____ (Eves)
May we publish the above information or a part thereof?_____
Comments_____

_____ I have enclosed $15.00 annual dues. _____ I have a packet station

Support

NAPRA
c/o Budd Churchward, WB7FHC
13304 131ST STREET KPN
Gig Harbor, Washington 98335

Support

LINK STATE AND AUDIO ALARM FOR TNC-1
USING PARALLEL USER PORT J-1

BY WA6OZJ, JIM & WD6FPY, BILL

REFERENCE PAGE 6-8 OF KIT MANUAL

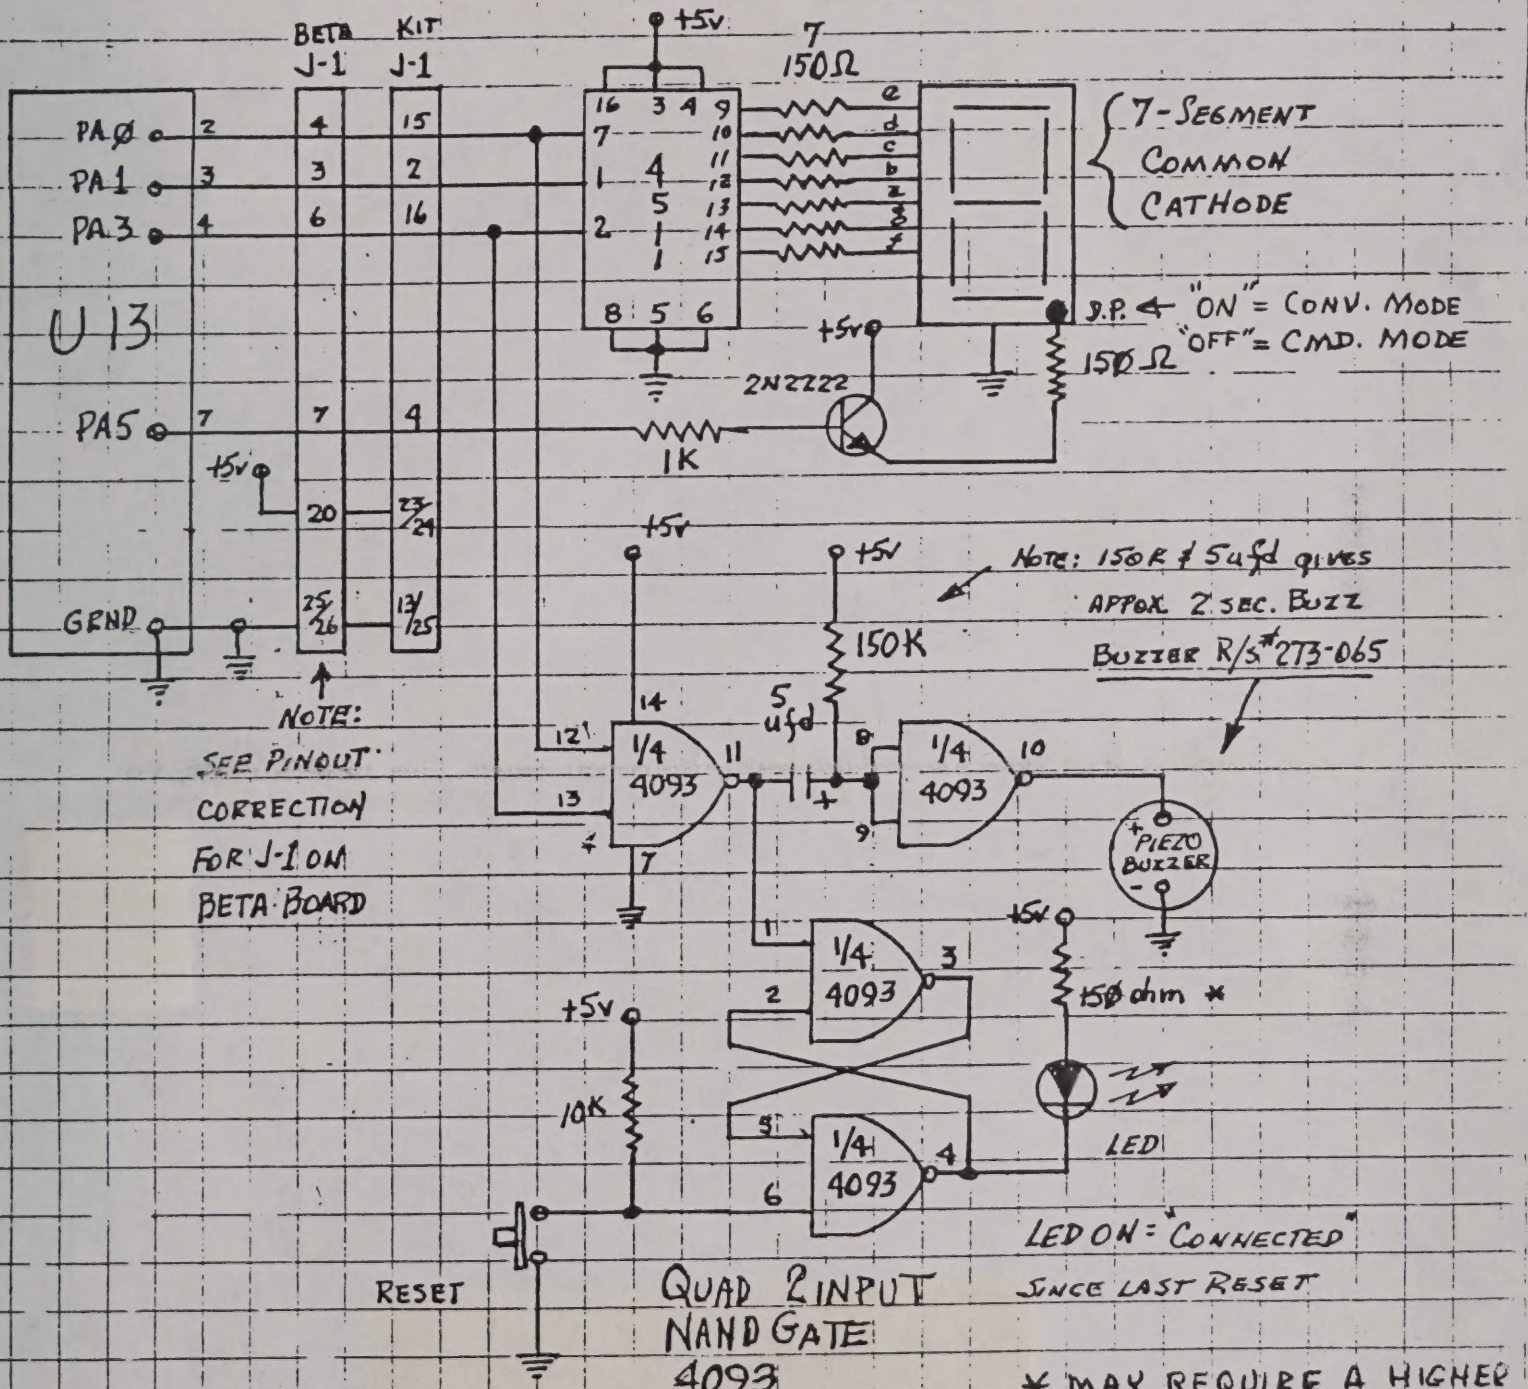

* MAY REQUIRE A HIGHER VALUE TO STAY LIT. TRY 470 Ω

-DE WA7HRA

IT IS TIME TO
RENEW YOUR NAPRA DUES

Northwest Amateur Packet Radio Association
13304 131st STREET KPN
Gig Harbor, WA 98335

KD7NM
Robert Donnell Jr.
P.O. Box 45645
Seattle WA 98145-0645